

Reducing State Explosion for Software Model Checking with Relaxed Memory Consistency Models

Tatsuya Abe¹, Tomoharu Ugawa², Toshiyuki Maeda¹, and Kousuke Matsumoto²

¹ {abet,tosh}@stair.center STAIR Lab, Chiba Institute of Technology

² {ugawa,matsumoto}@plas.info.kochi-tech.ac.jp Kochi University of Technology

Abstract. Software model checking suffers from the so-called state explosion problem, and relaxed memory consistency models even worsen this situation. What is worse, parameterizing model checking by memory consistency models, that is, to make the model checker as flexible as we can supply definitions of memory consistency models as an input, intensifies state explosion. This paper explores specific reasons for state explosion in model checking with multiple memory consistency models, provides some optimizations intended to mitigate the problem, and applies them to McSPIN, a model checker for memory consistency models that we are developing. The effects of the optimizations and the usefulness of McSPIN are demonstrated experimentally by verifying copying protocols of concurrent copying garbage collection algorithms. To the best of our knowledge, this is the first model checking of the concurrent copying protocols under relaxed memory consistency models.

Keywords: software model checking; relaxed memory consistency models; state explosion; reordering of instructions; integration of states; concurrent copying garbage collection

1 Introduction

Modern computing systems are based on concurrent/parallel processing designs for their performance advantages, and programs therefore must also be written to exploit these designs. However, writing such programs is quite difficult and error-prone, because humans cannot exhaustively consider the behaviors of computers very well. One approach to this problem is to use software model checking, in which all possible states that can be reached during a program's execution are explored. Many such model checkers have been developed (e.g., [12,18,26,25,7,8]).

However, most existing model checkers adopt *strict consistency* as a Memory Consistency Model (MCM) on shared memories, which only allows interleaving of instruction execution, and ignore more *relaxed* MCMs than strict consistency, which allow reorderings of instructions. This is not realistic because many modern computer architectures such as IA64, SPARC, and POWER [22,37,20] have adopted relaxed MCMs. Relaxed MCMs facilitate the performance of parallel-processing implementations because instructions may be reordered and multiple threads may observe distinct views on shared memory while strict consistency, which requires synchronization at each memory operation, is prohibitively expensive to be implemented on computer architectures.

As interest in MCMs has grown, some model checkers have introduced support for them [24,27,28,29]. However, these have been specific to certain MCMs, such as Total Store Ordering (TSO) and Partial Store Ordering (PSO) [10]. We are in the process of developing a model checker, McSPIN [9], that can handle multiple MCMs [1,2,3,5]. McSPIN can take an MCM as an input with a program to be verified. It has a specification language that covers various MCMs including TSO, PSO, Relaxed Memory Ordering (RMO), acquire and release consistency [23], Itanium MCM [21], and UPC MCM [39]. By using McSPIN, we can easily model check a *fixed* program under *various* MCMs.

However, software model checking suffers from the *state explosion problem*, and relaxed MCMs even worsen this, because the reordering of instructions allowed under relaxed MCMs enormously increases the number of reachable states. What is worse, parameterizing model checking by MCMs, that is, to make the model checker as flexible as we can supply definitions of MCMs as an input intensifies the state explosion.

This paper explains how model checking with multiple MCMs increases the number of reachable states, and clarifies the reasons for state explosion specific to model checking with multiple MCMs. In addition, some optimizations are provided that reduce state explosion, and their effects are demonstrated through experiments. The ideas behind the optimizations are simple: Pruning traces, partial order reduction, and predicate abstraction are well known to reduce state explosion in conventional model checking [17]. In our former paper [3], we arranged pruning traces and partial order reduction for model checking with relaxed MCMs. In this paper, we arrange predicate abstraction, and propose *stages*, which are integrations of states under relaxed MCMs.

Although the optimization in our earlier work have enabled verification of non-toy programs such as Dekker’s mutual exclusion algorithm [3], it was difficult to apply McSPIN to larger problems such as verifications of copying protocols of Concurrent Copying Garbage Collection algorithms (CCGCs), due to the state explosion. In this paper, we demonstrated the optimizations above enables McSPIN to verify larger programs; we checked if a desirable property of CCGCs, “in a single thread program, what the program reads is what it has most recently written”, are held or not for several CCGCs on multiple MCMs. Though we used verifications of GCs as examples in this paper, safety of GC is an important issue in the field (e.g., [14,16]), and this achievement is a positive development. To the best of our knowledge, this is the first model checking of copying protocols of CCGCs with relaxed MCMs.

The rest of this paper is organized as follows: Sect. 2 describes McSPIN with exploring the reasons for state explosion specific to model checking with MCMs, and Sect. 3 describes the relevant optimizations we have applied in McSPIN. Sect. 4 presents experimental results using McSPIN on different CCGCs and shows the effectiveness of the optimizations. Sect. 5 discusses related work, and the conclusions and directions for future work are presented in Sect. 6.

2 McSPIN

We first briefly review our earlier work [1,2,3] on constructing a general model checking framework with relaxed MCMs and developing and implementation. In the following,

we do not distinguish the framework from its implementation and refer to both as McSPIN. In McSPIN, threads on computers with shared memory are uniformly regarded as processes that have their own memories. Therefore, we formally call threads (in the usual sense) *processes* (in McSPIN), while we refer to them as “threads” when informally explaining behavior on shared-memory systems.

2.1 Syntax

A program is an N -tuple of sequences of instructions defined as follows:

(Instruction)	$i ::= \langle L, A, \iota \rangle,$
(Raw Instruction)	$\iota ::= \text{Move } r \ t \mid \text{Load } r \ x \mid \text{Store } x \ t \mid \text{Jump } L \text{ if } t \mid \text{Nop},$
(Term)	$t ::= v \mid r \mid t + t \mid t - t \mid \dots,$
(Attributes)	$A ::= \{a, \dots, a\},$

where N is the number of processes. An instruction i is a triple of a label, attributes, and a raw instruction. A label L designates an instruction in a program. An attribute $a \in A$ denotes an additional label for a raw instruction, has no effect itself, and are used to describe constraints specified by an MCM.

Here r is variable local to a process and x, y, \dots , are shared variables. The raw instruction $\text{Move } r \ t$ denotes the assignment of an evaluated value of a term t to a process-local variable r , which does not affect other processes. The term v denotes an immediate value. The terms $t_0 + t_1, t_0 - t_1, \dots$, denote standard arithmetic expressions. $\text{Load } r \ x$ represents loading x from its own memory and assigning its value to r . $\text{Store } x \ t$ denotes storing an evaluated value of t to x on its own memory. $\text{Jump } L \text{ if } t$ denotes a conditional jump to L depending on the evaluated value of t . Note that t contains no shared values; to jump to L depending on x , it is necessary to perform $\text{Load } r \ x$ in advance. Nop denotes the usual no-operation.

Careful readers may wonder why no synchronization instructions such as *memory fence* and *compare-and-swap* instructions appear. In McSPIN, a memory fence is represented as a Nop with attribute `fence`, and its effect is defined at each input MCM, that is, multiple types of fences can be defined. This flexibility enables verification of a *fixed* program with *different* MCMs as explained in more detail in Sec. 2.3. Compare-and-swap (usually an instruction on a computer architecture) is also represented by compound statements, which can be seen in Appendix A.

Programs (inputs to McSPIN) have to be written in the assembly-like modeling language. Such low-level languages are suitable for handling MCMs that require one to carefully take into account effects on specific computer architectures. However, these languages may not be practical for writing programs. McSPIN has a C-like modeling language to facilitate programming, but this is beyond the scope of the present paper.

2.2 Semantics

McSPIN adopts trace semantics with states. Execution traces are sequences of *operations*, defined as follows:

(Operation)	$o ::= \text{Fe}_q^j \ p \ i \mid \text{Is}_q^j \ p \ i \mid \text{Ex}_q^j \ p \ i \ \ell \ v \mid \text{Re}_q^j \ [p \Rightarrow p] \ i \ \ell \ v.$
-------------	---

One key point in handling different MCMs is to consider at most four kinds of operations for an instruction. For any instruction, its fetch and issue operations are considered. Load and store instructions have execution operations. Store instructions have reflect operations. An effect of each operation is formally defined in our former paper [5]. In this paper, we roughly explain why such operations are introduced.

Under very relaxed MCMs such as C++ [23] and UPC MCM [39], distinct threads can exhibit different program behaviors; that is, each thread has its own execution trace. To represent these in one trace, we add a process identifier q , denoting an observer process as a subscript of an operation. In addition, McSPIN can handle programs with loops. To distinguish multiple operations corresponding to an instruction, an operation has a branch counter j that designates the j th iteration within a loop.

We explain the four kinds of operations by example. $\text{Fe}_q^j p i$ denotes fetching an instruction i from a process p , which enables the issuance of i . By default, this also increments the program counter of p if i 's raw instruction is not `Jump`. If so, the program counter is not changed and will be changed when the `Jump` is issued. McSPIN is equipped with a *branch prediction* mode that can be switched on or off. In branch prediction mode, the program counter is non-deterministically incremented or set to L when `Jump L if t` is fetched. Thus, in order to handle branch prediction, fetch has to be distinguished from issue.

Although branch prediction is often ignored in specifications of MCMs, note that no branch prediction implicitly prohibits some kinds of reorderings across conditionals. For example, no branch prediction on the process-model that McSPIN adopts cannot perform the so-called *out-of-thin-air* read [30] in the program in Table 17.6 of Java language specification [32], although *legal* executions under Java MCM are specified by not using a *total* order of operations on such process-model but consistency between *partial* orders of operations on threads.

Operation $\text{Is}_q^j p i$ denotes the issuing of an instruction i to a process p . Effects that complete inside the register on p (not p 's own memory) are performed. For example, while issuing `Move r t` indicates assignment of an evaluated value of t to r , `Store x t` implies evaluation of t only. In branch prediction mode, a predicted execution trace in fetching `Jump` is checked.

Operation $\text{Ex}_q^j p i \ell v$ denotes execution of an instruction i on a process p . Effects that complete inside p are performed. For example, while the execution of `Load r x` means that v is loaded from x (at location ℓ) and assigned to r , `Store x t` represents storing an evaluated value v of t to x (at location ℓ) in p 's own memory. While an instruction is issued, its (intra-process) effect may not have occur yet. Itanium MCM allows such situation, by distinguishing issues from executions of instructions.

Operation $\text{Re}_q^j [p_0 \Rightarrow p_1] i \ell v$ denotes reflects of an instruction i from process p_0 to p_1 . The reflect of `Store x t` means storing an evaluated value v of t to x at ℓ in p_1 's own memory. While a store instruction is executed, that is, its effect is reflected to its store buffer, its (inter-process) effect may not be reflected to other processes yet. One reflect may be immediately passed, and another reflect may be delayed. Moreover, processes can observe distinct views a.k.a. *the IRIW test* [11]. Our definition covers such situations.

While the distinction enables delicate handling effects of instructions, it intensifies state explosion since the number of interleavings of operations increases.

To handle more relaxed MCMs, it is also necessary to distinguish multiple operations that are generated from an instruction in a loop statement, whereas this is unnecessary when queues can be used to handle specific MCMs such as TSO and PSO. For example, in a code (`Store x r0; Move r0 r0 + 1; Jump 0 if 1`) \parallel `Load r0 x`, the second fetch of the `Store` on the former process may follow the fetch of the `Load` on the latter process, while the first fetch of the `Store` on the former process may precede it. To the best of our knowledge, no existing method can handle such low-level jumps (across which instructions may be reordered) in a detailed fashion, which is necessary for verification of CCGCs.

2.3 Formalized Memory Consistency Models

MCMs are sets of constraints that control program behaviors on the very relaxed semantics that McSPIN adopts and are formally defined as a first-order formula as follows:

$$\varphi ::= x_\tau = x'_\tau \mid x_\tau < x'_\tau \mid \neg \varphi \mid \varphi \supset \varphi' \mid \forall x_\tau. \varphi(x_\tau),$$

where τ denotes one of Variable, Location, Label, Value, Instruction, Raw Instruction, Attribute, Branch Counter, and Operation. Here x_τ represents metavariables in the syntax of McSPIN. For example, $x_{\text{Location}} < x'_{\text{Location}}$ can be read as $\ell < \ell'$. In addition, $<$ with respect to Operation identifies the order of execution between operations. We use standard notation such as \wedge , \vee , and \exists and assign higher precedence to \neg , \wedge , \vee , and \supset .

Example constraints can be seen in [1,2,3], and Itanium and UPC MCMs are fully formalized in their journal version[5]. Here we focus on only two. In Sect. 2.1, we stated that the effect of a memory fence can be flexibly defined by an input MCM. A memory fence forces evaluation of all the reflects of store instructions that are fetched before the memory fence. This is represented as follows:

$$\text{Fe}_q^{j_0} p i_0 < \text{Fe}_q^{j_1} p (L_1, A_1, \text{Nop}) \supset \text{Re}_q^{j_0} [p \Rightarrow p_0] i_0 \ell_0 v_0 < \text{Is}_q^{j_1} p (L_1, A_1, \text{Nop}),$$

where $\text{fence} \in A_1$, i_0 's raw instruction is `Store`, and all free variables are universally quantified. Meanwhile, we can consider another operation that forces Load only:

$$\text{Fe}_q^{j_0} p i_0 < \text{Fe}_q^{j_1} p (L_1, A_1, \text{Nop}) \supset \text{Ex}_q^{j_0} p i_0 \ell_0 v_0 < \text{Is}_q^{j_1} p (L_1, A_1, \text{Nop}),$$

where $\text{fence} \in A_1$ and i_0 's raw instruction is `Load`.

One constraint that differentiates TSO from PSO with *multiple-copy-atomicity* [36], which prohibits two threads from observing different behaviors of write operations that the two threads do not perform, is whether reflects of store instructions are *atomically* performed *in program order*. This can be represented as follows:

$$\text{Fe}_q^{j_0} p i_0 < \text{Fe}_q^{j_1} p i_1 \supset \text{Re}_q^{j_0} [p \Rightarrow p_0] i_0 \ell_0 v_0 < \text{Re}_q^{j_1} [p \Rightarrow p_1] i_1 \ell_1 v_1,$$

where i_0 's and i_1 's raw instructions are `Store` instructions. This constraint causes *every* reflect of i_1 to await completion of *all* reflects of i_0 . Full constraints of TSO, PSO, and other relaxed MCMs are formalized in McSPIN's public repository [9].

2.4 Translation into PROMELA

McSPIN uses the model checker SPIN as an engine and translates programs written in our modeling language into PROMELA, the modeling language of SPIN. The underlying idea is quite simple. McSPIN translates sequential compositions of statements $i_0; i_1; \dots$ written in our modeling language into PROMELA loop statements as follows:

```

do
:: (guard0,0) -> (operation of Fe of  $i_0$ ); (epilogue0,0);
:: (guard0,1) -> (operation of Is of  $i_0$ ); (epilogue0,1);
:: (guard0,2) -> (operation of Ex of  $i_0$ ); (epilogue0,2);
:: (guard0,3) -> (operation of Re of  $i_0$  to  $p_0$ ); (epilogue0,3);
:: ...
:: (guard0,(N-1)+3) -> (operation of Re of  $i_0$  to  $p_{N-1}$ ); (epilogue0,(N-1)+3);
:: (guard1,0) -> (operation of Fe of  $i_1$ ); (epilogue1,0);
:: ...
:: else -> break;
od;

```

A PROMELA loop statement has multiple clauses with guards. One of those clauses whose guards are satisfied is non-deterministically chosen and processed. Let `clock` be a time counter. Each clause corresponds to performing an operation as follows:

```
end_o==0 -> o; end_o=clock; clock++;
```

where the positiveness of `end_o` denotes that o has already performed.

Although such a PROMELA code may admit very relaxed behavior that does not satisfy an input MCM, McSPIN appropriately removes such execution traces. Assertions can be written not only at the end of a program, but also at any place within. This is important for CCGC verification, because we would like to confirm data consistency at a certain place and moment. McSPIN modifies assertion statements to follow the input MCM. Let φ be an assertion that we wish to verify. McSPIN adds (formalized) constraints that an input MCM obligates to φ as a conjunct. For example, the constraint that differentiates TSO and PSO, as explained in Sect. 2.3, is translated into

$$!(\text{end_}\{\text{Fe}_q^{j_0} p i_0\} < \text{end_}\{\text{Fe}_q^{j_1} p i_1\}) \mid \mid \text{end_}\{\text{Re}_q^{j_0} [p \Rightarrow p_0] i_0 \ell_0 v_0\} < \text{end_}\{\text{Re}_q^{j_1} [p \Rightarrow p_1] i_1 \ell_1 v_1\}$$

and added to the assertion φ as a conjunct, where `!` and `||` represent negation and disjunction in PROMELA, respectively. Thus, execution traces that violate the MCM are removed when assertions are checked.

3 Optimizations

Here we provide MCM-sensitive optimization techniques to reduce the problem specific to model checking with multiple MCMs. The optimizations described in Sects. 3.1 and 3.2 were introduced in [3]; we briefly review them here in order to make it easy to understand an optimization introduced in Sect. 3.3.

3.1 Enhanced Guards: Pruning Inadmissible Execution Traces

As explained in Sect. 2.4, McSPIN explores all execution traces and removes traces that are inadmissible under an input MCM in checking assertions. This is obviously redundant. A straightforward method to prune inadmissible execution traces is to enhance guards for clauses corresponding to operations. A guard that is uniformly generated as $\text{end_o} == 0$ from an operation o in Sect. 2.4 is enhanced by an input MCM (details are provided in [3]). We explain this using the constraint that differentiates TSO and PSO, as set out in Sect. 2.3. The constraint claims that all reflects of i_1 must wait for all reflects of i_0 , where i_0 precedes i_1 in program order. McSPIN adds a condition

$$! (\text{end_}\{\text{Fe}_q^{j_0} p i_0\} < \text{end_}\{\text{Fe}_q^{j_1} p i_1\}) \mid \text{end_}\{\text{Re}_q^{j_0} [p \Rightarrow p_0] i_0 \ell_0 v_0\} > 0$$

corresponding to this claim to the guard of the reflect of i_1 .

3.2 Defining Predicates: Promoting Partial Order Reduction

As explained in Sect. 2.4, it is necessary to judge whether an execution trace is admissible to a given MCM. This means that it is also necessary to remember orders between operations in the execution trace. The most straightforward method is to use a time counter; that is, to substitute a variable end_o (defined at each operation) with the time at which operation o was performed. However, time counters are too concrete to reduce state explosion. For example, consider four operations o_0, o_1, o_2, o_3 under the constraint $o_0 < o_1 \supset o_2 < o_3$. If times are substituted for the variables end_o_k ($0 \leq k < 4$), then the number of combinations $\langle \text{end_o}_0, \text{end_o}_1, \text{end_o}_2, \text{end_o}_3 \rangle$ is 24 ($=4!$), which distinguishes states more concretely than the constraint requires.

When considering the constraint rule, it suffices to remember the order of o_0 and o_1 and of o_2 and o_3 , because nothing else is used to define the constraint. We introduce new variables $\text{ord_o}_0\text{-o}_1$ and $\text{ord_o}_2\text{-o}_3$, and call them *defining predicates* of the constraint or, formally, atomic formulas consisting of the predicate symbol $<$ (or \leq) between operations that occur in the constraint. Because the defining predicates preserve the order of times at which the operations are performed, we change end_o_k to boolean variables that denote whether the operation has been performed. After all the operations have been performed (that is, $\text{end_o}_k = 1$ ($0 \leq k < 4$)), the possible states are $\langle \text{ord_o}_0\text{-o}_1, \text{ord_o}_2\text{-o}_3 \rangle = \{ \langle 0, 0 \rangle, \langle 0, 1 \rangle, \langle 1, 1 \rangle \}$, of cardinality 3.

3.3 Stage: Abstracting Programs by MCM-Deriving Predicates

Predicate abstraction [17] is one promising method to reduce state explosion in model checking. In this subsection, we show that predicates exist that are determined by an input MCM. Such predicates integrate states that do not have to be separated with respect to an input MCM. Therefore, the predicate abstractions have no omission of checking.

To handle the effects of instructions more delicately, McSPIN has at most four kinds of operations, Fe, Is, Ex, and Re for one instruction. However, some MCMs do not require complete distinction. Assume that an input MCM has the constraint $\text{Is}_q^j p i < o \supset \text{Ex}_q^j p i \ell_1 v_1 < o$ as called *integration* in [2,3], which indicates that no operation can interleave two operations $\text{Is}_q^j p i$ and $\text{Ex}_q^j p i \ell_1 v_1$. In an earlier version, McSPIN

generated clauses that had guards waiting for $\text{Ex}_q^j p i \ell_1 v_1$ when $\text{Is}_q^j p i$ was performed. Such guards control program behaviors in accordance with an input MCM.

In this paper, we promote integration to state level rather than execution-trace level. In earlier versions, McSPIN generated one clause at each operation; that is, at most $3 + N$, the cardinality of $\{\text{Fe}, \text{Is}, \text{Ex}\} \cup \{\text{Re } k \mid 0 \leq k < N\}$, clauses at each instruction, where N is again the number of processes, and $\text{Re } k$ denotes a reflect to k . In the current version, McSPIN can accept additional input *stages* $S = \{s_0, s_1, \dots, s_{M-1}\}$ for an input MCM. Formally, stages are partitions of $\{\text{Fe}, \text{Is}, \text{Ex}\} \cup \{\text{Re } k \mid 0 \leq k < N\}$. We write f_S for the induced mapping from the stages. McSPIN generates M clauses at each instruction i , where M is the number of stages of i as follows:

```

do
:: (guards0, fS(s0)) -> (operation of fS(s0) of i0); (epilogue0, fS(s0));
:: (guards0, fS(s1)) -> (operation of fS(s1) of i0); (epilogue0, fS(s1));
:: ...
:: (guards0, fS(sM-1)) -> (operation of fS(sM-1) of i0 to pN-1); (epilogue0, fS(sM-1));
:: ...
:: else -> break;
od;

```

This optimized translation reduces checking space and time. By loading such a PROMELA code, SPIN remembers not unintegrated states themselves but stages. This implies that state-vector on SPIN is kept small. Memory is not, therefore, consumed so much. This optimization also saves time to check whether clauses are executable since the number of clauses is smaller.

Let us see example stages for TSO and PSO with neither *branch prediction* nor *multiple-copy-atomicity* [36], which prohibits two threads from observing different behaviors of write operations that the two threads do not perform. Since these MCMs allow Loads to overtake (inter-process) effects of Stores, each member of $\{\text{Re } k \mid 0 \leq k < N\}$ has to be separated from Ex . However, Fe , Is , and Ex do not have to be separated. Also, $\text{Re } k$ does not have to be distinguished from $\text{Re } k'$ ($k' \neq k$) by multiple-copy-atomicity. We can therefore introduce the following stages:

$$S = \{s_0, s_1\} \quad f_S(s) = \begin{cases} \{\text{Fe}, \text{Is}, \text{Ex}\} & \text{if } s = s_0 \\ \{\text{Re } k \mid 0 \leq k < N\} & \text{if } s = s_1 \end{cases}.$$

Given a stage S (and its mapping f_S), McSPIN automatically returns PROMELA code in which clauses are integrated; in particular, guards and epilogues are appropriately generated from an input MCM.

4 Experiments

In this section, we demonstrate the effects of the optimizations introduced in Sect. 3. The figure to the right shows our experimental environment, with ample memory.

The optimizations described in Sects. 3.1 and 3.2 have enabled verification of relatively large programs such as Dekker's algorithm [3]. Here we demonstrate that the optimization described in Sect. 3.3 enables verification of genuinely large programs.

CPU:	Intel Xeon E5-2670 2.6GHz
Memory:	DDR3-1066 1.5TB
SPIN:	6.4.5
GCC:	5.3.0

4.1 Experimental Setting

We chose CCGCs as examples of large programs. In this subsection, we briefly explain the CCGCs we used.

Garbage collection (GC) is a basic service of modern programming languages. Its role is to find garbage, that is, data objects that are no longer in use by the application, and to reclaim the memory that those objects occupy. Copying GC accomplishes this by copying live objects, i.e., those that may be used in the future, to a separate space and then releasing the old space that contains the copied objects and garbage. Concurrent GC, as the name suggests, runs concurrently with the application. What is difficult in designing CCGC algorithms is that the garbage collector thread and an application thread may race; the application thread may change the contents of an object that is being copied by the garbage collector. This may be the case even with a single thread application. Because an application thread changes, or mutates, the object, we call it a *mutator*. If a mutator writes to the object that is being copied, the collector may copy a stale value, which means that the latest value gets lost. Various copying protocols have been proposed to provide application programmers with reasonable MCMs, all of which require the mutators to do some work on every read (*read barrier*) or write (*write barrier*) operation or both, in which the mutator synchronizes with the collector.

Because such barriers incur overhead for every read or write operation, one goal of CCGC algorithms is to design barriers that are as lightweight as possible. Thus, synchronizations such as compare-and-swap should be minimized. With relaxed MCMs, memory barriers should also be minimized. Unfortunately, the synchronizations required for safety depend on the given MCM; it is often the case that those synchronizations that are redundant for one MCM are mandatory for another.

Model. We experimentally checked the safety of concurrent copying protocols, *in a single thread program, what the mutator reads is what it has most recently written*. This property is expected to be held in any reasonable MCMs such as the happens-before consistency of Java [32]. The complete McSPIN models for checking this property can be found in Appendix or the McSPIN public repository [9]. Here, we briefly explain the model.

In our model checking, we made some assumptions. We assume that there is a single mutator thread, i.e., the application is a single thread program. Remark that even if there is a single mutator thread, there is another thread, the collector thread, and they may race. We also assume that there is only a single object with a single integer slot in the heap.

The mutator has a pointer to the object and repeatedly reads from and writes to the object through the pointer. On write operations, it remembers the value it wrote. On read operations, it checks if the read value is equal to the value it lastly wrote. Meanwhile, the collector copies the object following to the copying protocol of each algorithm. Once it successfully copied, the collector rewrites the mutator's pointer to the object so that the pointer points at the copy.

To cooperate with the collector, the mutator uses the read and write barriers required by the copying protocol on its read and write operations. For some algorithms, the mutator also performs so called the checkpoint operation between object accessing

operations, where the mutator polls and answers collector’s requests. Some collectors request the mutator to answer the handshake by setting a per-mutator handshake request flag. The checkpoint operation clears the flag to let the collector know the mutator has observed the flag set. In TSO, if a mutator observes the flag is set, all stores preceding the store setting the flag are guaranteed to be visible to the mutator.

We created McSPIN models for each CCGC algorithms we describe below. In the models, the mutator has an infinite loop, where it reads or writes once per an iteration. It also performs a checkpoint operation before and after each read or write. Thus, the supremum of loop iterations on the mutator limits the number of mutator’s memory accesses.

GC Algorithm In this paper, we checked three GC algorithms: Chicken [34], Staccato [31] and Stopless [33]. The details of these algorithms can be found in their papers. Here, we briefly explain their features.

Chicken and Staccato were basically the same algorithm though they are developed independently. The only difference is their target MCMs; Chicken is designed for the MCMs of Intel CPU such as IA64 [22], while Staccato’s main target seems to be POWER MCM [20]. These algorithms use compare-and-swap operations to resolve races between the collector and a mutator. In the IA64 MCM, the compare-and-swap is usually realized by the instruction sequence `lock cmpxchg`. This sequence implies memory fences. As for POWER MCM, the manual [20] shows a sample implementation of the compare-and-swap operation that does not imply memory fences.

Stopless is a different algorithm from those two. It uses compare-and-swap operations that implies memory fence excessively, hence chances of reordering are fewer.

4.2 Effect of Optimization

In this subsection, we reveal the effectness of the *stage* optimization described in Sect. 3.3. We verified the models created in Sect. 4.1 by using McSPIN with and without the optimization. We fixed the supremum of iteration on the collector to 1 and varied that on the mutator from 1 to 2.

Tab. 1 shows the results of the verification. Note that any PROMELA code produced by McSPIN consumed around 170 MB of memory as constant overhead. As Tab. 1 shows, the amounts of memory consumed and elapsed times are greatly reduced in all algorithms compared with those without the optimization. In particular, when the supremum of loop iterations on the mutator was set to 2, McSPIN without the optimization often required around 1 TB of memory, which is far from reasonable. However, a single iteration could not detect any error even for the algorithms that actually work incorrectly with PSO, i.e., Chicken and Stopless.

Tab. 1 also suggests that the more instructions the model had, the more effective the stage optimization was. For example, in Chicken and Staccato with a single iteration, the net memory consumption was reduced to 3.9–8.8 %, while, in Stopless, it was reduced to 3.3 and 4.4 %. This is because the stage optimization reduced the number of units that are subject to reordering, or clauses of the do-loop in the PROMELA code.

To the best of our knowledge, this is the first model checking of these algorithms with PSO, due to the optimizations given in this paper.

Table 1. Effects of optimization: In TSO, a compare-and-swap instruction implies memory fences. In PSO, it does not. Columns labeled with “col.” and “mut.” list the number of instructions of the collector and the mutator, respectively. Column labeled with “loop” lists the supremum of loop iterations on the mutator. For verification either with or without stages, the first column shows the results; ✓ means no error was found and × means a violation was found. The following columns are the number of state transitions, the amount of memory consumed, and elapsed times for verification, respectively. Columns labeled with “mem. ratio” and “time ratio” list the ratios of memory and time consumption for verification with/without stages. The column labeled with “net mem. ratio” lists those that do not count constant overhead.

MCM	algorithm	col.	mut.	loop	without stages			with stages			mem. ratio (%)	net mem. ratio (%)	time ratio (%)		
					state (K)	memory (MB)	time (sec.)	state (K)	memory (MB)	time (sec.)					
TSO	chicken	24	42	1	✓	108	8,595	132	✓	23	908	8	10.6	8.8	6.2
				2	✓	2,506	546,038	8,637	✓	534	24,960	433	4.6	4.5	5.0
	staccato	32	46	1	✓	141	14,918	236	✓	26	1,032	11	6.9	5.8	4.7
				2	✓	3,888	1,097,491	16,022	✓	733	43,432	735	4.0	3.9	4.6
	stopless	33	87	1	✓	90	28,183	378	✓	14	1,404	19	5.0	4.4	5.0
				2	✓	564	564,635	7,705	✓	87	18,885	585	3.3	3.3	7.6
PSO	chicken	24	42	1	✓	308	25,208	430	✓	65	1,652	28	6.6	5.9	6.6
				2	×	1,136	264,857	4,248	×	237	12,190	227	4.6	4.5	5.3
	staccato	32	46	1	✓	143	15,166	243	✓	26	1,032	11	6.8	5.7	4.7
				2	✓	4,020	1,144,602	16,975	✓	751	44,920	768	3.9	3.9	4.5
	stopless	33	87	1	✓	177	55,210	833	✓	30	2,520	46	4.6	4.3	5.5
				2	×	45	45,416	630	×	8	2,148	41	4.7	4.4	6.6

Table 2. Variants of Staccato

algorithm	col.	mut.	loop	TSO			PSO				
				state (K)	memory (MB)	time (sec.)	state (K)	memory (MB)	time (sec.)		
staccato	32	46	1	✓	26	1,032	11	✓	26	1,032	11
			2	✓	733	43,432	735	✓	751	44,920	768
staccato_pso	31	44	1	✓	25	908	10	✓	25	1,032	10
			2	✓	719	38,721	637	✓	755	40,953	703
staccato_bug	30	44	1	✓	28	1,032	11	✓	35	1,156	14
			2	✓	819	43,184	726	×	235	12,810	217

4.3 Reducing Memory Fences of Staccato

Because Staccato is designed for the relaxed MCM of POWER, some memory fences are redundant on a stricter MCM. Thus, we designed and verified a variant of Staccato for a PSO MCM with a compare-and-swap that does not imply memory fences. In addition, we created an incorrect variant that lacks mandatory fences for the PSO MCM. These variants are labeled `staccato_pso` and `staccato_bug`.

The result of verification is shown in Tab. 2. The verification is conducted with the stage optimization. The result of `staccato_bug` shows that McSPIN detected an error if we reduced fences too much.

The variants of Staccato demonstrate the usefulness of McSPIN. When we modify a GC algorithm for a machine with some MCM that is different from the one that the GC is originally designed for, we add or remove some synchronizations. However, the modified model often lacks synchronizations. McSPIN can detect such errors in the variant with a reasonable memory consumption. This enables us to check the GC when we are performing modifications.

<pre> active proctype main() { run mem(); run proc0(); ... } proctype proc0() { WRITE(x0,0); WRITE(x1,1); } ... </pre>	<pre> proctype mem() { endmem: do ::atomic{COMMIT_WRITE(queue_proc0);} ::atomic{COMMIT_WRITE(queue_proc1);} ... od; } inline WRITE (var,val) {...} inline COMMIT_WRITE (queue) {...} </pre>
--	---

Fig. 1. Hand-written code**Table 3.** Comparison between McSPIN and hand-coding

	TSO						PSO					
	McSPIN			hand-written			McSPIN			hand-written		
	state	memory (MB)	time (sec.)	state	memory (MB)	time (sec.)	state	memory (MB)	time (sec.)	state	memory (MB)	time (sec.)
1	25	0.006	0.01	19	0.002	0.01	25	0.006	0.02	19	0.003	0.01
2	52	0.017	0.02	60	0.009	0.00	65	0.021	0.02	79	0.017	0.01
3	116	0.053	0.02	149	0.026	0.00	241	0.110	0.05	337	0.095	0.01
4	241	0.153	0.04	313	0.064	0.02	977	0.619	0.20	1,405	0.504	0.02
5	457	0.391	0.08	585	0.143	0.01	3,985	3.405	0.98	5,749	2.500	0.07
6	800	0.897	0.18	1,004	0.276	0.01	16,145	18.107	6.13	23,269	11.894	0.31
7	1,312	1.882	0.35	1,615	0.493	0.01	65,041	93.290	34.06	93,637	54.294	1.66
8	2,041	3.659	0.61	2,469	0.829	0.02	261,137	468.195	171.10	375,685	246.497	8.28

4.4 McSPIN vs. Hand-Coding

In this subsection, we compare PROMELA codes generated by McSPIN with codes written by hand and confirm how close McSPIN is to an ideal implementation.

Whereas McSPIN generates uniform PROMELA codes that contain variables to remember orders between operations, etc., to support different MCMs, some variables are essentially unnecessary for verifications specific to TSO and PSO. Because TSO never reorders store instructions, queues (for all shared variables) at each thread to buffer effects of write instructions suffice for verifications under TSO as shown in Fig. 1. The two `WRITE`s put `<x0, 0>` and `<x1, 1>` into the queue in order. Reads from the queue to shared memory are performed by `COMMIT_WRITE`s on a process `mem`. We omit the implementation details. For PSO, one queue at *each* shared variable is enough to reorder the effects of write instructions to distinct shared variables.

Tab. 3 compares PROMELA codes generated by McSPIN with those written by hand where the constant overhead is removed. The programs are simple, consisting of multiple store instructions (without loops). Verified properties are fixed to be true. Each column is similar to Tab. 1. The digits in the names of the codes denote the number of store instructions at each thread, respectively. The number of states almost coincides. Slight differences appear to derive from the current implementation of SPIN, because we observe that SPIN returns fewer states for a PROMELA code with a loop statement and control variables (such as code generated by McSPIN) than another PROMELA code with a sequential composition of statements (like hand-written code). However, we have not investigated this in detail.

McSPIN consumes more memory and time. This is a result of the sizes of the state vectors and is inevitable, because McSPIN defines more variables to determine program structures than hand-written codes, as explained in the beginning of this subsection.

5 Related Work

There exists no work, which is directly compared with our work, of model checking to take multiple MCMs in a uniform way. Therefore, we can find no work for its optimization has been studied.

Jonsson's seminal work discovered the potential of SPIN for program translation toward model checking with relaxed MCMs [24]. However, he could not conduct a large number of experiments, because his program translation was not completely automatic and optimized. This paper has addressed the problems that he left open. McSPIN supports various MCMs and takes an MCM as an input, and its program translation is automatic. McSPIN is greatly optimized and enables verification of larger concurrent algorithms such as copying protocols of CCGCs.

Linden et al. [27,28,29] tackled the state explosion problem by representing store buffers as automata. However, they handled relatively strict relaxed MCMs such as TSO and PSO, unlike McSPIN. It is an open issue to extend their representation so as to handle more relaxed MCMs and apply it to McSPIN.

Modex [19], a model extractor of SPIN that is guided by a user-defined test-harness, translates C codes into PROMELA codes. However, Modex ignores relaxed MCMs. Although revising Modex so as to handle relaxed MCMs is surely one approach, we have developed McSPIN in order to show the potential of program translation toward model checking with relaxed MCMs with no restriction derived from the existing tool.

Travkin et al. [40] developed a similar tool that translates programs into PROMELA codes and uses SPIN as the engine for model checking, demonstrated verifications of linearizability of concurrent algorithms under TSO, and planned to tackle PSO. However, their translator, which generates codes that are similar to hand-written PROMELA code as introduced in Sect. 4.4, cannot be immediately applied to relaxed MCMs beyond PSO. Unlike their approach, ours supports relaxed MCMs by virtue of constructing a base that allows such relaxed behaviors and then defines MCMs as constraints on the base. Although an issue of our approach is addressing the state explosion problem, this paper has presented optimizations for the problem.

Dan et al. [13] reported high utility of predicate abstractions in model checking with relaxed MCMs by verifying some programs with predicate abstractions under TSO and PSO. They proposed the notion of predicate *extrapolation* to abstract a boolean program for an input program. Although the stages introduced in this paper can be regarded as predicate abstractions, there is a difference in usage: McSPIN considers at most four kinds of multiple states at one instruction to support various MCMs beyond TSO/PSO. Although it is necessary to handle the worst case under the most relaxed MCM, this is not always the case. Stages are states that are integrated by predicates that are uniformly generated by an input MCM. Therefore, abstractions by the predicates never leak out of checking. Dan et al.'s technique of extrapolating predicates seems to be compatible with stages, and its combination with stages is an open issue.

Theorem proving in program logic is also one promising approach to program verification with relaxed MCMs [35,15,41,4,6]. Formal verifications of GC algorithms with relaxed MCMs using theorem provers have recently appeared [16]. However, fully automated verification by model checking is usually preferable to manual (or semi-automatic) construction of proofs in theorem proving.

6 Conclusion and Future Work

We have explained the reasons for the state explosion problem specific to model checking with multiple MCMs, presented optimizations modified from pruning execution traces, partial order reduction, and predicate abstraction, and applied them to McSPIN, our model checker with MCMs. We have also shown the effectiveness of the optimizations through experiments of verifications of copying protocols of CCGCs, which are larger programs.

There are four future directions for this work. Although we verified copying protocols of CCGCs as examples of large programs in this paper, a verification of GC algorithm is itself subject of our interest. Our future work includes verifications of wide range of GC algorithms and other properties such as wait-freedom for Chicken, which the authors designed as a wait-free CCGC [34]. These verifications may require more complicated settings including pointers and/or multiple mutators, which need still larger models. The second is to show a verification of concurrent copying protocols with MCMs that are more relaxed than PSO. An advantage of McSPIN is its ability to support various MCMs. The third is to show more realistic benchmark programs, e.g., SV-COMP benchmarks [38]. The fourth is further optimization of McSPIN to verify even larger programs.

Acknowledgments The authors thank the anonymous reviewers for several comments to improve the final version of the paper. This research partly used computational resources under Collaborative Research Program for Young Scientists provided by Academic Center for Computing and Media Studies, Kyoto University. This work was supported by JSPS KAKENHI Grant Numbers 25871113, 25330080, and 16K21335.

References

1. Abe, T., Maeda, T.: Model checking with user-definable memory consistency models. In: Proc. of PGAS, short paper. (2013) 225–230
2. Abe, T., Maeda, T.: A general model checking framework for various memory consistency models. In: Proc. of HIPS. (2014) 332–341
3. Abe, T., Maeda, T.: Optimization of a general model checking framework for various memory consistency models. In: Proc. of PGAS. (2014)
4. Abe, T., Maeda, T.: Concurrent program logic for relaxed memory consistency models with dependencies across loop iterations. Journal of Information Processing (2016) To appear.
5. Abe, T., Maeda, T.: A general model checking framework for various memory consistency models. International Journal on Software Tools for Technology Transfer (2016) To appear.
6. Abe, T., Maeda, T.: Observation-based concurrent program logic for relaxed memory consistency models. In: Proc. of APLAS. (2016) To appear.
7. Abe, T., Maeda, T., Sato, M.: Model checking with user-definable abstraction for partitioned global address space languages. In: Proc. of PGAS. (2012)
8. Abe, T., Maeda, T., Sato, M.: Model checking stencil computations written in a partitioned global address space language. In: Proc. of HIPS. (2013) 365–374
9. Abe, T., Maeda, T., Ugawa, T.: McSPIN. <https://bitbucket.org/abet/mcspin/>.
10. Adve, S., Gharachorloo, K.: Shared memory consistency models: a tutorial. Computer 29(12) (1996) 66–76

11. Boehm, H.J., Adve, S.V.: Foundations of the C++ concurrency memory model. In: Proc. of PLDI. (2008) 68–78
12. Cavada, R., Cimatti, A., Jochim, C.A., Olivetti, G.K.E., Pistore, M., Roveri, M., Tchaltev, A.: NuSMV User Manual. 2.5 edn. (2002)
13. Dan, A., Meshman, Y., Vechev, M., Yahav, E.: Predicate abstraction for relaxed memory models. In: Proc. of SAS. (2013) 84–104
14. Doligez, D., Gonthier, G.: Portable, unobtrusive garbage collection for multiprocessor systems. In: Proc. of POPL. (1994) 70–83
15. Ferreira, R., Feng, X., Shao, Z.: Parameterized memory models and concurrent separation logic. In: Proc. of ESOP. (2010) 267–286
16. Gammie, P., Hosking, T., Engelhardt, K.: Relaxing safely: Verified on-the-fly garbage collection for x86-TSO. In: Proc. of PLDI. (2015) 99–109
17. Graf, S., Saidi, H.: Construction of abstract state graphs with PVS. In: Proc. of CAV. (1997) 72–83
18. Holzmann, G.J.: The SPIN Model Checker. Addison-Wesley (2003)
19. Holzmann, G.J., Smith, M.H.: An automated verification method for distributed systems software based on model extraction. *IEEE Trans. Software Eng* **28**(4) (2002) 364–377
20. IBM: PowerPC Architecture Book, Version 2.02. (2005)
21. Intel: A Formal Specification of Intel Itanium Processor Family Memory Ordering. (2002)
22. Intel: Intel 64 and IA-32 Architectures Software Developer’s Manual. (2016)
23. ISO/IEC 14882:2011: Programming Language C++. (2011)
24. Jonsson, B.: State-space exploration for concurrent algorithms under weak memory orderings: (preliminary version). *SIGARCH Computer Architecture News* **36**(5) (2008) 65–71
25. Kroening, D., Tautschnig, M.: CBMC - C bounded model checker - (competition contribution). In: Proc. of TACAS. Volume 8413 of LNCS. (2014) 389–391
26. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Proc. of CAV. Volume 6806 of LNCS. (2011) 585–591
27. Linden, A., Wolper, P.: An automata-based symbolic approach for verifying programs on relaxed memory models. In: Proc. of SPIN. Volume 6349 of LNCS. (2010) 212–226
28. Linden, A., Wolper, P.: A verification-based approach to memory fence insertion in relaxed memory systems. In: Proc. of SPIN. Volume 6823 of LNCS. (2011) 144–160
29. Linden, A., Wolper, P.: A verification-based approach to memory fence insertion in PSO memory systems. In: Proc. of TACAS. Volume 7795 of LNCS. (2013) 339–353
30. Manson, J., Pugh, W., Adve, S.V.: The Java memory model. In: Proc. of POPL. (2005) 378–391
31. McCloskey, B., Bacon, D.F., Cheng, P., Grove, D.: Staccato: A parallel and concurrent real-time compacting garbage collector for multiprocessors. Report RC24504, IBM (2008)
32. Oracle: The Java Language Specification. Java SE 8 edition edn. (2015)
33. Pizlo, F., Frampton, D., Petrank, E., Steensgaard, B.: Stopless: a real-time garbage collector for multiprocessors. In: Proc. of ISMM. (2007) 159–172
34. Pizlo, F., Petrank, E., Steensgaard, B.: A study of concurrent real-time garbage collectors. In: Proc. of PLDI. (2008) 33–44
35. Ridge, T.: A rely-guarantee proof system for x86-TSO. In: Proc. of VSTTE. (2010) 55–70
36. Sarkar, S., Sewell, P., Alglave, J., Maranget, L., Williams, D.: Understanding POWER multiprocessors. In: Proc. of PLDI. (2011) 175–186
37. SPARC International, Inc.: The SPARC Architecture Manual, Version 9. (1994)
38. SV-COMP: Competition on Software Verification. <https://sv-comp.sosy-lab.org/>.
39. The UPC Consortium: UPC Language Specifications Version 1.3. (2013)
40. Travkin, O., Mütze, A., Wehrheim, H.: SPIN as a linearizability checker under weak memory models. In: Proc. of Haifa Verification Conference. Volume 8244 of LNCS. (2013) 311–326
41. Vafeiadis, V., Narayan, C.: Relaxed separation logic: A program logic for C11 concurrency. In: Proc. of OOPSLA. (2013) 867–884

A McSPIN Model – Compare-and-Swap

Here we show the McSPIN model of compare-and-swap instructions, which we used in the models of each copying protocols of CCGCs. We made four variants; CAS compare-and-swaps a single variable, and CAS2 does two variables at the same time, and those that do not report if the operation succeeded or not, whose names have the `_NORET` suffix. CAS2 is used for modeling a double-word compare-and-swap and bit fields. The fences (`McSPIN_fence`) appear for TSO but not for PSO.

```

/* compare-and-swap */
static inline CAS(int target, int old, int new, int retval) {
#pragma McSPIN attribute atomic
{
    if (target == old) {
        target = new;
        retval = 1;
    } else
        retval = 0;
}
}

/* two-variable compare-and-swap */
static inline CAS2(int target1, int target2, int old1, int old2, int new1, int new2,
int retval) {
#pragma McSPIN attribute atomic
{
    if (target1 == old1)
        if (target2 == old2) {
            target1 = new1;
            target2 = new2;
            retval = 1;
        } else
            retval = 0;
    else
        retval = 0;
}
}

/* compare-and-swap */
static inline CAS_NORET(int target, int old, int new) {
#pragma McSPIN attribute atomic
{
    if (target == old)
        target = new;
}
}

/* two-variable compare-and-swap */
static inline CAS2_NORET(int target1, int target2, int old1, int old2, int new1, int new2) {
#pragma McSPIN attribute atomic
{
    if (target1 == old1)
        if (target2 == old2) {
            target1 = new1;
            target2 = new2;
        }
}
}

```

B McSPIN Model for Chicken

Here we show our McSPIN model for the copying protocol of Chicken.


```

#include "stdbool.h"
#include "atomic.h"

#define INITIAL_VALUE 0

#define FROM_OBJECT 0
#define TO_OBJECT 1

int main()
{
    /* handshake */
    int hs_req;
    /* heap */
    int from_header_fwd, from_header_copying, from_body;
    int to_body;
#define to_header_fwd TO_OBJECT
#define to_header_copying FALSE
    /* mutator's variable */
    int root;

#pragma McSPIN parallel sections
    {
        /* collector */
#pragma McSPIN section
        {
            int success = 0;
            while (success == 0) {
                from_header_copying = TRUE;

                start_handshake();
                wait_for_handshake();

                /* assume: to_header_copying = FALSE */
                /* assume: to_header_fwd = TO_OBJECT */
                to_body = from_body;

                CAS2(from_header_fwd, from_header_copying, FROM_OBJECT, TRUE, TO_OBJECT, FALSE,
                    success);
            }
            root = from_header_fwd; /* flip */
        }

        /* mutator */
#pragma McSPIN section
        {
            int last_written = INITIAL_VALUE;
            int readval;
            while (true) {
#pragma McSPIN nondeterministic
                {
                    ack_handshake();
                }
#pragma McSPIN nondeterministic
                {
                    last_written = 1 - last_written;
                    write(root, last_written);
                } else {
                    read(root, readval);
                    McSPIN_assert(McSPIN_variable(last_written,1,0) == McSPIN_variable(readval,1,0));
                }
            }
#pragma McSPIN nondeterministic
            {
                ack_handshake();
            }
        }
    }
}

```

```

}

static inline ack_handshake()
{
    if (hs_req == true)
        hs_req = false;
}

static inline start_handshake()
{
    hs_req = true;
}

static inline wait_for_handshake()
{
    while (hs_req != 0)
        ;
}

/* read barrier */
static inline read(int obj, int retval) {
    if (obj == FROM_OBJECT) {
        if (from_header_fwd == FROM_OBJECT)
            retval = from_body;
        else
            retval = to_body;
    } else {
        /* assume: to_header_fwd = TO_OBJECT */
        retval = to_body;
    }
}

/* write barrier */
static inline write(int obj, int val) {
    if (obj == FROM_OBJECT) {
        if (from_header_copying == TRUE) {
            CAS2_NORET(from_header_fwd, from_header_copying, FROM_OBJECT, TRUE, FROM_OBJECT, FALSE);
        }
        if (from_header_fwd == FROM_OBJECT)
            from_body = val;
        else
            to_body = val;
    } else {
        /* assume: to_header_copying = FALSE */
        /* assume: to_header_fwd = TO_OBJECT */
        to_body = val;
    }
}

```

C McSPIN Model for Staccato

Here we show our McSPIN model for the copying protocol of Staccato. Two fences are removed for `staccato_pso` and another fence is omitted for `staccato_bug`. They are marked with comments in the following model.

```

#include "stdbool.h"
#include "atomic.h"

#define INITIAL_VALUE 0

#define FROM_OBJECT 0
#define TO_OBJECT 1

```

```

int main()
{
    /* handshake */
    int hs_req;
    /* heap */
    int from_header_fwd, from_header_copying, from_body;
    int to_body;
    #define to_header_fwd TO_OBJECT
    #define to_header_copying FALSE
    /* mutator's variable */
    int root;

    #pragma McSPIN parallel sections
    {
        /* collector */
        #pragma McSPIN section
        {
            int success = 0;
            while (success == 0) {
                from_header_copying = TRUE;
                McSPIN_fence();

                start_handshake();
                wait_for_handshake();
                McSPIN_fence(); /* removed for staccato_pso and saccato_bug */

                /* assume: to_header_copying = FALSE */
                /* assume: to_header_fwd = TO_OBJECT */
                to_body = from_body;

                McSPIN_fence(); /* omitted for saccato_bug */
                start_handshake();
                wait_for_handshake();

                CAS2(from_header_fwd, from_header_copying, FROM_OBJECT, TRUE, TO_OBJECT, FALSE,
                    success);
            }
            root = from_header_fwd; /* flip */
        }

        /* mutator */
        #pragma McSPIN section
        {
            int last_written = INITIAL_VALUE;
            int readval;

            while (true) {
                #pragma McSPIN nondeterministic
                {
                    ack_handshake();
                }
                #pragma McSPIN nondeterministic
                {
                    last_written = 1 - last_written;
                    write(root, last_written);
                } else {
                    read(root, readval);
                    McSPIN_assert(McSPIN_variable(last_written,1,0) == McSPIN_variable(readval,1,0));
                }
            }
            #pragma McSPIN nondeterministic
            {
                ack_handshake();
            }
        }
    }
}

```

```

static inline ack_handshake()
{
    McSPIN_fence();
    if (hs_req == true)
        hs_req = false;
    McSPIN_fence(); /* removed for staccato_pso and saccato_bug */
}

static inline start_handshake()
{
    hs_req = true;
}

static inline wait_for_handshake()
{
    while (hs_req != 0)
        ;
}

/* read barrier */
static inline read(int obj, int retval) {
    if (obj == FROM_OBJECT) {
        if (from_header_fwd == FROM_OBJECT)
            retval = from_body;
        else
            retval = to_body;
    } else {
        /* assume: to_header_fwd = TO_OBJECT */
        retval = to_body;
    }
}

/* write barrier */
static inline write(int obj, int val) {
    if (obj == FROM_OBJECT) {
        if (from_header_copying == TRUE) {
            CAS2_NORET(from_header_fwd, from_header_copying, FROM_OBJECT, TRUE, FROM_OBJECT, FALSE);
        }
        if (from_header_fwd == FROM_OBJECT)
            from_body = val;
        else
            to_body = val;
    } else {
        /* assume: to_header_copying = FALSE */
        /* assume: to_header_fwd = TO_OBJECT */
        to_body = val;
    }
}

```

D McSPIN Model for Stopless

Here we show our McSPIN model for the copying protocol of Stopless.

```

#include "stdbool.h"
#include "atomic.h"

/* status */
#define IN_ORIGINAL 0
#define IN_WIDE 1
#define IN_COPY 2

/* value */
#define INITIAL_VALUE 0

```

```

/* space */
#define FROM_OBJECT 0
#define WIDE_OBJECT 1
#define TO_OBJECT 2

int main()
{
    /* from object */
    int from_fwd = FROM_OBJECT;
    int from_data = INITIAL_VALUE;

    /* wide object */
    int wide_status = IN_ORIGINAL;
    int wide_data;

    /* to object */
#define to_fwd TO_OBJECT
    int to_data;

    /* root */
    int root = FROM_OBJECT;

#pragma McSPIN parallel sections
    {
        /* collector */
#pragma McSPIN section
        {
            int x;
            int success;

            /* allocate wide object and install its forwarding pointer */
            CAS_NORET(from_fwd, FROM_OBJECT, WIDE_OBJECT);
            /* copy the payload to the wide object */
            x = wide_data;
            CAS2_NORET(wide_status, wide_data, IN_ORIGINAL, x,
                       IN_WIDE, from_data);
            /* copy the payload to the final copy */
            success = 0;
            while (!success) {
                x = wide_data;
                to_data = x;
                CAS2(wide_status, wide_data, IN_WIDE, x, IN_COPY,
                    x, success);
            }
            /* repoint to the final copy */
            from_fwd = TO_OBJECT;

            root = from_fwd; /* flip */
        }

        /* mutator */
#pragma McSPIN section
        {
            int last_written = INITIAL_VALUE;
            int readval;
            int write_tmp;
            int write_success;

            while (true) {
#pragma McSPIN nondeterministic
                {
                    read(root, readval);
                    McSPIN_assert(McSPIN_variable(last_written,1,0) == McSPIN_variable(readval,1,0));
                } else {
                    last_written = 1 - last_written;
                    write(root, last_written);
                }
            }
        }
    }
}

```

```

    }
  }
}

/* read barrier */
static inline read(int obj, int retval) {
  if (obj == FROM_OBJECT) {
    if (from_fwd == FROM_OBJECT)
      retval = from_data;
    else if (from_fwd == WIDE_OBJECT) {
      if (wide_status == IN_ORIGINAL)
        retval = from_data;
      else if (wide_status == IN_WIDE)
        retval = wide_data;
      else if (wide_status == IN_COPY)
        retval = to_data;
    } else if (from_fwd == TO_OBJECT)
      retval = to_data;
  } else { /* obj == TO_OBJECT */
    /* assume: to_fwd == TO_OBJECT */
    retval = to_data;
  }
}

/* write barrier */
static inline write(int obj, int val) {
  if (obj == FROM_OBJECT) {
    if (from_fwd == FROM_OBJECT)
      CAS_NORET(from_fwd, FROM_OBJECT, WIDE_OBJECT);
    if (from_fwd == WIDE_OBJECT) {
      write_success = 0;
      while(!write_success) {
        if (wide_status == IN_ORIGINAL) {
          write_tmp = wide_data;
          CAS2(wide_data, wide_status, write_tmp, IN_ORIGINAL, val, IN_WIDE, write_success);
        } else if (wide_status == IN_WIDE) {
          write_tmp = wide_data;
          CAS2(wide_data, wide_status, write_tmp, IN_WIDE, val, IN_WIDE, write_success);
        } else if (wide_status == IN_COPY) {
          to_data = val;
          write_success = 1;
        }
      }
    } else if (from_fwd == TO_OBJECT)
      to_data = val;
  } else { /* obj == TO_OBJECT */
    /* assume: to_fwd == TO_OBJECT */
    to_data = val;
  }
}
}

```